RUTGERS–NEW BRUNSWICK
Institute for Teaching, Innovation, and Inclusive Pedagogy

Designing AI-Aware Assignments

A resource for Rutgers University Instructors

Overview of this Resource

Generative AI (GenAI) can now complete many of the tasks instructors have traditionally asked students to perform on their own.

The technologies can draft essays, solve problem sets, summarize readings, write code, and much more, often in seconds. For instructors who design assignments, this raises an important question. **Instead of asking only “What do I want students to do?” instructors also have to ask “What do I want students to learn by doing it?”**

If you have found yourself wondering whether an assignment you have relied on for years still accomplishes what you wanted it to do, you are not alone. The assignment itself has not changed, but students now have access to tools that can bypass some of the thinking it was designed to develop. It is easy to feel caught between assigning it as you always have or abandoning it altogether. This resource starts from a different premise: there is usually another option.

Rather than beginning with AI tools or policies, this resource begins with assignment design. It starts by asking what students are meant to learn, then uses that answer to decide what role, if any, generative AI should play. **This is AI-aware assignment design (Vee, Watkins, & Bruff, 2026).**

An AI-aware assignment reflects the instructor making an intentional decision about the role GenAI should play, grounded in the assignment’s learning outcomes and communicated clearly to students. That decision might be to use GenAI, leave it out, or something in between. An instructor who decides that generative AI has no place in a particular assignment has done just as much AI-aware design as one who incorporates it throughout. What matters is that the choice is intentional and connected to student learning.

Goals of this Resource

Every assignment you give now carries a question it did not carry several years ago: What role should generative AI play in this one? This resource helps you answer it, task by task.

Specifically, it will help you:

- think deliberately about the role generative AI should (or should not) play in your assignments;
- keep student learning at the center of decisions about generative AI,
- communicate your expectations clearly so students understand when and why generative AI is (or is not) appropriate; and
- connect with Rutgers resources that can support your work.

This resource is not a mandate to adopt generative AI, nor is it a checklist to complete. The ideas and examples throughout are offered as approaches to consider, not prescriptions to follow. You know your discipline, your students, and your learning outcomes best. The goal of the resource is simply to offer a framework for making intentional, well-informed decisions about AI-aware assignment design.

Resource Layout

This resource has two sections (see Figure 1). You do not need to read this resource from beginning to end. While each section can stand on its own, **The Design Method** is the foundation for everything that follows and the best place to begin.

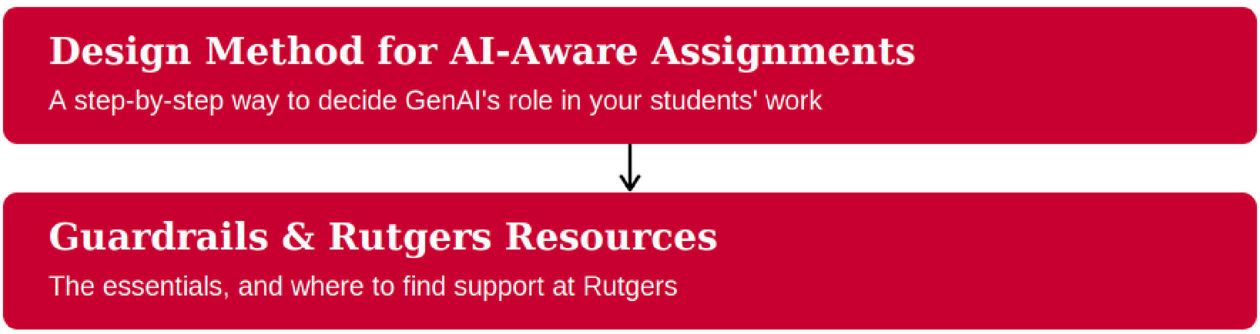

Figure 1. The two sections of this resource

Design Method for AI-Aware Assignments

Designing an AI-aware assignment begins with a simple question: **What role, if any, should generative AI play in this assignment?** Rather than answering that question once for an entire course, this resource encourages instructors to answer it task by task, where the most important design decisions are made.

The Design Method consists of four steps:

1. **Start with your learning outcomes, using backward design (Wiggins & McTighe, 2005).**
2. **Find out what generative AI can do with each task.**
3. **Think through a decision, with the help of Maha Bali's cake-making analogy.**
4. **Name and communicate that decision to students, one option is to use the AI Assessment Scale.**

The steps work the same way whatever you end up deciding, whether a task calls for no GenAI at all or invites it in as a full partner.

Step 1: Start with the course learning outcomes

Backward design (Wiggins & McTighe, 2005) reminds us to begin not with the assignment (or with generative AI), **but with the learning itself**.

Rather than treating an assignment as a single unit, break it into the individual tasks students complete, and ask what they are meant to learn from each one.

This first step matters because it sets the standard that every later decision is measured against.

If you are not sure what a task is for, you cannot really judge whether GenAI strengthens it or weakens it. Breaking the assignment into parts, rather than looking at it as a whole, is what lets us make our decisions.

A simple way into this step is to finish the sentence "By doing this, students will be able to..." for each part of the assignment. If you find you cannot finish it, that task may be worth reconsidering, regardless of GenAI.

💡 **For example: A data analysis assignment is rarely just "report the results." It is several learning tasks layered together: cleaning and preparing the data, choosing an appropriate method, running the analysis, and interpreting what the findings mean. Each step might call for a different decision about the usage of generative AI.**

Step 2: Stress-test the task against GenAI

Once you know what each task is for, the next step is to find out what generative AI can already do with it. Try the task in a current generative AI tool using the kind of prompt a student might reasonably write.

This move is similar to the DEER framework for AI-aware assignment design, which asks instructors to define an assignment's purpose, evaluate it by testing the task with GenAI, explore how others in the discipline are handling it, and reflect on the results (Cummings et al., 2024 as cited in Vee at al., 2026).

This is not about catching students or making the assignment "AI-proof," which is rarely possible with today's tools. It is about gathering information.

Seeing what GenAI can and cannot do helps us make informed design decisions rather than relying on assumptions.

Sometimes you might discover that GenAI produces a polished version of exactly the thinking you wanted students to do for themselves, which is a signal that the task needs our attention.

💡 **For example: Ask generative AI to write the reflection essay you assign after an experiential-learning placement. You may find it produces something fluent but generic, lacking the specific nuance that comes from a student's actual hours in the field. That tells you the task is more resilient, and that your criteria should reward exactly the lived experience GenAI cannot invent.**

Step 3: Use the cake analogy to decide the role of GenAI

Now you can decide, task by task, what role GenAI should play, and here Maha Bali's cake-making analogy offers a helpful way to think it through (see Figure 2).

It sorts tasks by how much cognitive work they ask of the student and how much GenAI assistance is appropriate. The analogy also mirrors the progression from building foundational knowledge toward higher-order thinking described in Bloom's taxonomy (Anderson & Krathwohl, 2001; Bloom et al., 1956).

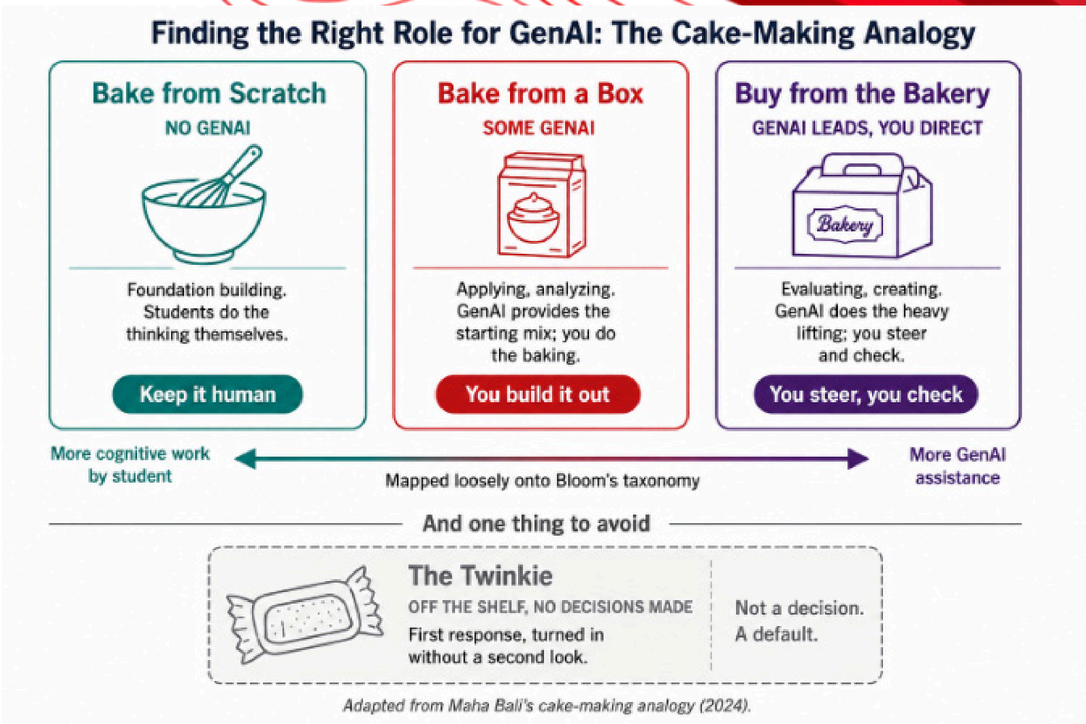

Figure 2. The cake-making analogy for GenAI decisions. Adapted from Maha Bali's cake-making analogy, 2024. Image created through ChatGPT Instant 5.5 (OpenAI), July 2026.

Some tasks call for students to **bake from scratch**, with no GenAI at all. These are our foundation-building assignments, where students learn the facts, concepts, and procedures they will need later.

They need to do this work themselves because it is hard to think critically about something you do not yet understand. You learn to read music before you can critique a performance. When a task is building that foundation, GenAI bypasses the very learning the task exists to create, and the decision is to keep it in the student's hands.

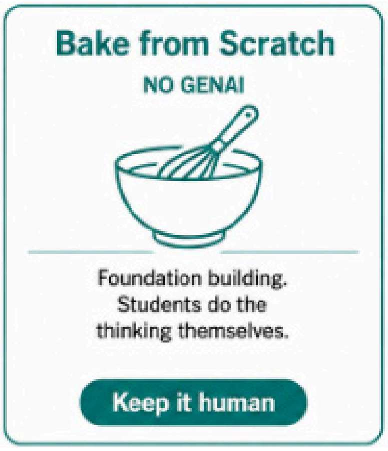

Other tasks invite students to **bake from a box mix**, which is to say, GenAI provides some structure, an outline or a different definition for example, and students do the real work of developing it into their own.

Here, students know enough to start applying and analyzing, and GenAI can act as a thought partner that helps them brainstorm, explore "what-if" scenarios, or notice connections they might have missed. Students are still doing the thinking; GenAI is supporting it. The decision in these cases is to permit GenAI in a defined, supporting role.

Sometimes it makes sense to **buy the cake from the bakery**, working closely with GenAI to complete the task, directing it and shaping its output the way you would work with a professional baker to get the cake you want. This works because the student has already developed real expertise. After all, you only catch the misspelled name on the cake if you already know how it should read.

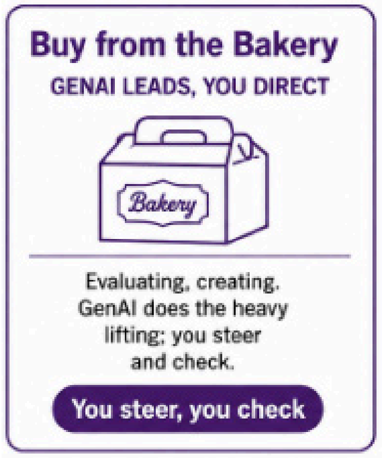

At these higher Bloom's levels, such as evaluating and creating, the skill is no longer in performing every step yourself but in knowing what good work looks like and how to get there and directing GenAI to reach it. The decision here is to allow GenAI to lead, paired with the student's active prompting and final check.

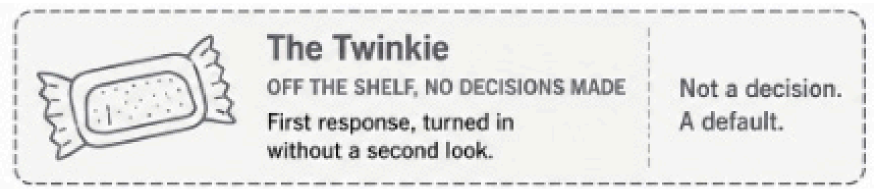

And then there is the final option: skipping the kitchen and the bakery altogether and **grabbing something off the shelf at the store, a Twinkie**. It looks like a cake, and it will work at the last minute, but nobody involved made any decisions about it. It is the same no matter who buys it or what the occasion calls for.

That is what happens when you take GenAI's first response and turn it in without looking at it twice: no direction, no fit to the task, no one home. It's not that GenAI did too much of the work here, the bakery tier does that too, it's that no judgment touched the result at any point.

💡 **For example: In an introductory programming course, writing basic loops and functions by hand is bake from scratch, since this is where students build the mental model they will need later. Using GenAI to generate a starter outline for a project, which students then develop and complete themselves is baking from a box mix. Directing GenAI to generate boilerplate setup code in an advanced course, where students already know what good code looks like and check it carefully, is buy from the bakery. Copying whatever GenAI produces for a reflective assignment without reading it over is the Twinkie, and it is what this whole method is built to help you avoid.**

Thinking through the analogy keeps our attention where it belongs, on how much cognitive work you want the student to do, and how actively they are directing GenAI rather than simply accepting what it returns. If a task is strengthened by involving GenAI, you build GenAI in on purpose. If GenAI might diminish the learning, that task is worth keeping as a human-only one.

As Ludwig and Zakrajsek (2026) indicate, the aim is to amplify student thinking rather than outsource it, and "understanding when not to use GenAI is as important as knowing when to use it" (p. 89).

This is also why choosing not to use generative AI needs no special defense. Deciding that a foundational task should be done from scratch is not a rejection of GenAI. It is backward design doing exactly what it is supposed to do.

Step 4: Be transparent with students about how GenAI can be used

An AI-aware assignment only works if students understand the decisions behind it. The AI Assessment Scale (AIAS) gives a shared vocabulary for the what (see Figure 3).

Where Step 3 helped with thinking through a decision, the AIAS gives that decision a clear, named place you can point students to. The thinking still comes first, and the scale simply names where it lands. The AIAS is one of many tools that can support how you and your students can engage with GenAI in your assignments. There are other frameworks like the PAIRR (Peer and AI Review + Reflection, Sperber, 2025) and the Online Learning Consortium's GenAI Use and Ethics Framework that you could also find useful.

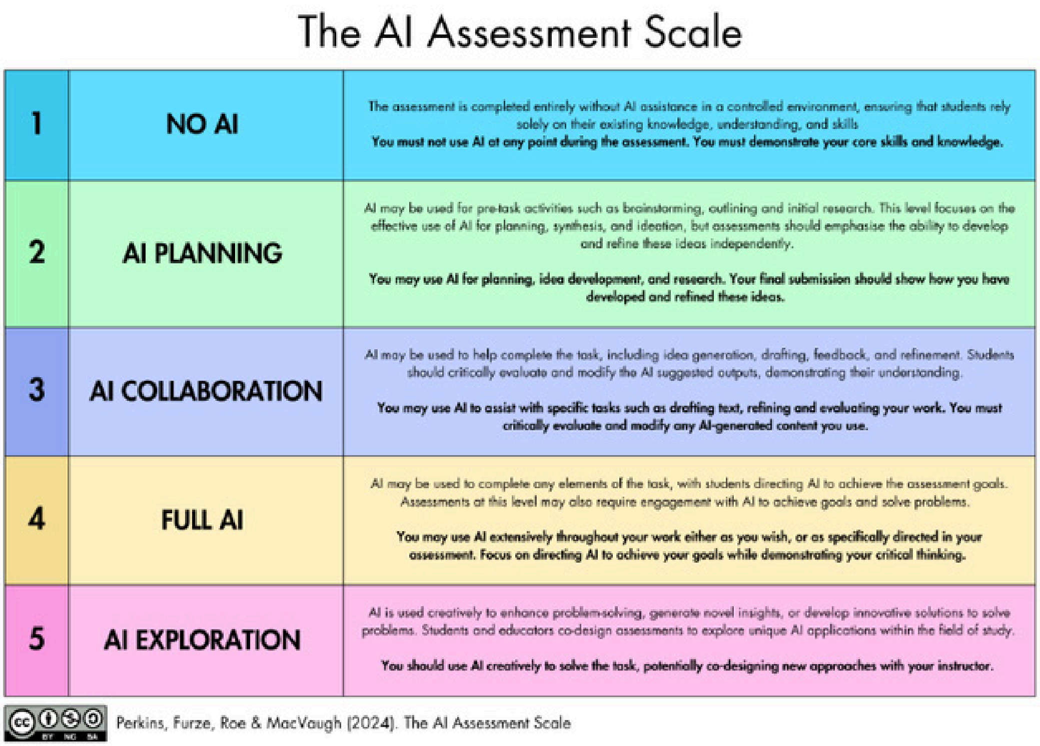

The AI Assessment Scale

1	NO AI	The assessment is completed entirely without AI assistance in a controlled environment, ensuring that students rely solely on their existing knowledge, understanding, and skills **You must not use AI at any point during the assessment. You must demonstrate your core skills and knowledge.**
2	AI PLANNING	AI may be used for pre-task activities such as brainstorming, outlining and initial research. This level focuses on the effective use of AI for planning, synthesis, and ideation, but assessments should emphasise the ability to develop and refine these ideas independently. **You may use AI for planning, idea development, and research. Your final submission should show how you have developed and refined these ideas.**
3	AI COLLABORATION	AI may be used to help complete the task, including idea generation, drafting, feedback, and refinement. Students should critically evaluate and modify the AI suggested outputs, demonstrating their understanding. **You may use AI to assist with specific tasks such as drafting text, refining and evaluating your work. You must critically evaluate and modify any AI-generated content you use.**
4	FULL AI	AI may be used to complete any elements of the task, with students directing AI to achieve the assessment goals. Assessments at this level may also require engagement with AI to achieve goals and solve problems. **You may use AI extensively throughout your work either as you wish, or as specifically directed in your assessment. Focus on directing AI to achieve your goals while demonstrating your critical thinking.**
5	AI EXPLORATION	AI is used creatively to enhance problem-solving, generate novel insights, or develop innovative solutions to solve problems. Students and educators co-design assessments to explore unique AI applications within the field of study. **You should use AI creatively to solve the task, potentially co-designing new approaches with your instructor.**

Figure 3. The AI Assessment Scale from Perkins, Furze, Roe & MacVaugh (revised 2025)

The five AIAS levels run from no GenAI use through to full, GenAI-led exploration, which gives a place to locate any task you have thought through.

Used this way, the scale is a tool for communicating, not a shortcut around the thinking you did in Step 3. You are thinking about your learning outcomes and then telling students, in language instructors all share, where the task landed, rather than picking a level off the shelf and working backward from it.

It helps to communicate the why right alongside the what. This transparency draws on the Transparency in Learning and Teaching (TILT) framework, which asks instructors to make the purpose, tasks, and criteria of an assignment explicit to students (Winkelmes, 2013). When you tell students that a task is meant to be done from scratch, with no GenAI, because this is where they build the foundation they will rely on all semester, you are doing more than setting a rule.

You are teaching the very judgment instructors hope they will develop, which is a sense of when a tool supports their thinking and when it gets in the way of it.

Transparency runs both ways. Just as you tell students which AIAS level applies to a task, you can ask them to be open about how they used GenAI to complete it, whether that is a sentence describing what a tool helped with, a note on which parts are their own work, or a short reflection on what the tool did and did not do for their thinking.

Whatever the task or the AIAS level, the same underlying test applies: can the student evaluate what GenAI gave them, verify what needed verifying, make the work genuinely their own, and explain and defend what they submit?

That reflection is where much of the learning lives, since it asks students to notice when GenAI supported their thinking and when it stood in for it. A simple disclosure line, paired with the AIAS level you have set, makes students' choices visible and gives us a clearer picture of their learning. Your syllabus is a good place to set specific expectations.

💡 For example: On a research paper prompt, a single line tells students where each task landed and why: "Developing your argument is yours alone (AIAS Level 1), because this is the thinking the paper exists to build. You may use AI to check your citation formatting (AIAS Level 2), as long as you verify every entry yourself." Then you might ask students for a matching line back: a short note where students say how they used GenAI, if they did, and what it did or did not do for their work.

One Assignment, start to finish: A literature review

Literature reviews are everywhere in higher education, and they are among the assignments most disrupted by GenAI, which makes them a useful case to walk through.

Step 1: Start with the outcome. Broken into tasks (locating sources, judging credibility, synthesizing across studies, identifying a gap), a literature review is really several learning tasks at once.

Step 2: Stress-test the assignment. GenAI can produce a fluent synthesis in seconds, the very thinking you most want students to do, so that is where the risk concentrates.

Step 3: Decide how much GenAI to use. Synthesis stays human (bake from scratch); generating search terms allows GenAI as a partner (bake from a box mix); reference formatting can be outsourced, then checked (buy from the bakery).

Step 4: Make the GenAI decision transparent to students. Each decision goes to students as an AIAS level with its reason, so one assignment ends up spanning three positions, which is the method working exactly as intended.

Notice what happened here. The literature review was never really lost. Taken apart, it turned out that the thinking you cared about most, the judgment and the synthesis, was exactly where GenAI is weakest. **The assignment did not have to go. It had to be seen clearly.** These same four steps serve all of us, whatever you end up deciding.

AI-Free by Design

Sometimes, the most pedagogically sound decision is not to use generative AI at all.

When the learning outcomes depend on students doing the cognitive work themselves, an AI-free assignment isn't a rejection of technology. Instead, it reflects an intentional design choice grounded in the same principles that guide every other AI-aware assignment.

It is easy to assume that "AI-aware" means "AI-using." An instructor who deliberately decides that generative AI has no place in an assignment has still designed that assignment with AI in mind. As Ludwig & Zakrajsek (2026) remind us, "understanding when not to use GenAI is as important as knowing when to use it" (p. 89). A course built mostly or entirely around GenAI-free assignments can be thoroughly AI-aware.

The important distinction is between **ignoring** generative AI and **making an intentional decision** about it. One leaves students guessing your expectations; the other clearly communicates how the assignment supports learning and why GenAI is or is not appropriate.

💡 **For example: You notice that a written reflection assigned after class was too easy to complete with GenAI's help, and too hard to tell apart from a student's own thinking once it was. Rather than trying to police the writing, you change the format: students still write something, but they also talk it through aloud with a peer, and someone captures what actually got said. The learning you care about, honest self-assessment, still happens. It just happens somewhere GenAI cannot easily reach.**

💡 **Designing GenAI-free assignments well: TIIP's resource on Alternative Assessment & Grading offers a range of approaches that fit naturally here, including oral exams, two-stage exams, and authentic or experiential tasks that keep students actively engaged in the thinking the assignment is designed to develop.**

That is The Design Method: a way to decide, task by task, what role GenAI should play in your students' work, whether that role is large, small, or none at all. Whatever you decide, a few things are worth keeping in mind, no matter how you land. The next section covers the essentials.

Guardrails and Rutgers Resources

Whatever you decide about generative AI, a few things remain true. This closing section highlights a few guardrails to keep in mind and points to Rutgers resources that can support your work.

Three guardrails to keep in mind:

1. **Lead with design, not detection.** It is tempting to reach for GenAI detection tools, but they are unreliable and can be biased, flagging a student's original work as machine-written, sometimes at high rates and disproportionately for multilingual writers (Liang et al., 2023).

As Ellis and Lodge (2024) argue, the more productive move is to stop looking for evidence of cheating and start looking for evidence of learning. A better approach is to design assignments with clear, purposeful expectations about generative AI that are tied to what students are meant to learn. The Design Method offers one way to do that exactly. A well-designed assignment does more to support integrity than any detector.

2. **Clearly indicate the role of GenAI, and say why.** This builds on the transparency step of The Design Method, carried up to the level of the whole syllabus. Students should be able to tell, for any given assignment, what role GenAI may play and the reasoning behind that choice. Clear, assignment-specific language, ranging from "no AI at any point during this assessment" to "use AI creatively to solve this task," removes guesswork and models the kind of decision-making instructors hope students will develop. If your expectations differ from one assignment to the next, which can happen given the process in The Design Method section, make sure students are also aware of this.

3. **Don't require tools some students cannot access.** If an assignment depends on a particular GenAI tool, access becomes a fairness question. Some tools carry a subscription cost, and students who can pay for a more capable version may produce stronger work than those who cannot. Rutgers provides several GenAI tools at no additional cost to all students, faculty, and staff, so when GenAI is part of required work, it is worth pointing students to those.

4. **Where you can, give students a choice**. When GenAI is permitted rather than required, students can decide whether to use it or not within the parameters you set. That serves both students who want practical experience with these tools and students who would rather not use them. Choice works best with structure: clear parameters, an alternative option with the same learning outcomes, and a short reflection on how they approached the assignment. TIIP's Student Choice Resource Library resource offers more information, including implementation recommendations and sample activities.

Rutgers Resources

From Rutgers IT

The Artificial Intelligence at Rutgers page is the place to start, and it is kept current as tools and policies change. A few things you will find there:

- **GenAI tools at no additional cost**, available to students, faculty, and staff: Microsoft Copilot Chat, Gemini Notebook (formerly NotebookLM), and Google Gemini.
- Data privacy guidance, including a data classification chart for AI tools that shows what kinds of information are and are not appropriate to enter into GenAI tools. Protected or identifiable student information should never be entered into a GenAI tool that has not been approved. See also Information Classification Policy 70.1.2.

- University policies relevant to GenAI use, including Academic Integrity 10.2.13, Information Classification 70.1.2, and Acceptable Use of IT Resources 70.1.1.
- General AI guidance and governance, including the AI@Rutgers Initiative steering committee and Rutgers' guidance on the use of AI.
- AI literacy resources, including a growing library of LinkedIn Learning courses on generative AI, covering prompt design, ethical implications, and practical uses.

From the Institute for Teaching, Innovation, and Inclusive Pedagogy (TIIP)

TIIP offers ongoing spaces to think through this work together. This resource grew out of those conversations, and we would love to have you join us.

- **Teaching and GenAI Pathways Program**. A self-paced program where you can work through milestones in generative AI teaching practices. Earn digital badges along the way.
- **GenAI Community of Practice**. A monthly space to think through live questions about generative AI with colleagues, grounded in real classroom practice. Come to share what you are trying, or just to listen and hear what others are doing.
- **Additional Emerging Technologies Resources**: The Institute for Teaching, Innovation, and Inclusive Pedagogy has developed and curated other resources, including advice for students.

From Other Campus Units

Support for teaching with GenAI comes from many different spaces across Rutgers. A few worth knowing about:

- Teaching Critical AI Literacies (SAS, Critical AI). Critical AI, a Rutgers-based interdisciplinary journal and initiative, offers critical perspectives on generative AI along with peer-reviewed classroom resources for teachers and students.

- AI in Teaching Libguide (Rutgers University Libraries). Includes sample syllabus language, plus some examples of AI literacy assignments.
- Artificial Intelligence Libguides (Rutgers University Libraries). Start here for a grounding in generative AI terminology, plus how to cite it.
- UOES Resource Library (University Online Education Services). Practical guides for online and hybrid courses, including assignment instructions when using GenAI
- Critical Approaches to AI (Rutgers University Libraries). A look at the ethical, environmental, and labor concerns surrounding generative AI.

Acknowledgements

Thank you to the TIIP team and colleagues in TIIP's GenAI Community of Practice who read early drafts and offered feedback that shaped this resource.

We want to especially thank **Tatiana Rodriguez** (School of Communication and Information; School of Management and Labor Relations) for her contribution: catching a gap in the cake-making analogy, offering a sharper way to think about student ownership of GenAI-assisted work, and contributing the classroom example behind the AI-Free by Design section.

GenAI Disclosure Statement

Claude (Anthropic) was used in developing this resource to assist with structure, research summarization, drafting, and graphics. All framing, pedagogical decisions, and final content were determined and reviewed by the author. This disclosure models the practice the resource describes: considering generative AI's role deliberately and being transparent about the choice.

References

Anderson, L. W., & Krathwohl, D. R. (2001). A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives. Longman.

Bali, M. (2024). Cake-Making Analogy for Setting Generative AI Guidelines/Ethics. In Buyserie, B. & Thurston, T. N. (Ed.), Teaching and Generative AI: Pedagogical Possibilities and Productive Tensions. Utah State University. https://doi.org/10.26079/e204-acc5

Bloom, B. S., Englehart, M. D., Furst, E. J., Hill, W. H., & Krathwohl, D. R. (1956). Taxonomy of educational objectives: The classification of educational goals. Handbook I: Cognitive domain. David McKay.

Cummings, R. E., Monroe, S. M., & Watkins, M. (2024). Generative AI in first-year writing: An early analysis of affordances, limitations, and a framework for the future. Computers and Composition, 71, 102827. https://doi.org/10.1016/j.compcom.2024.102827

Ellis, C., & Lodge, J. M. (2024, July 8). Stop looking for evidence of cheating with AI and start looking for evidence of learning. LinkedIn.

Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. Patterns, 4(7), 100779. https://doi.org/10.1016/j.patter.2023.100779

Ludwig, L., & Zakrajsek, T. (2026). The Science of Learning Meets AI. Routledge. https://doi.org/10.4324/9781003689102

Perkins, M., Furze, L., Roe, J., & MacVaugh, J. (2024). The Artificial Intelligence Assessment Scale (AIAS): A framework for ethical integration of generative AI in educational assessment. Journal of University Teaching and Learning Practice, 21(6), 49-66. https://doi.org/10.53761/q3azde36

Sperber, L., MacArthur, M., Minnillo, S., Stillman, N., & Whithaus, C. (2025). Peer and AI Review+ Reflection (PAIRR): A human-centered approach to formative assessment. Computers and Composition, 76, 102921. https://doi.org/10.1016/j.compcom.2025.102921

Vee, A., Watkins, M., & Bruff, D. (2026). The Norton Guide to AI-Aware Teaching. W. W. Norton. https://wwnorton.com/books/9781324127673

Wiggins, G., & McTighe, J. (2005). Understanding by Design (2nd ed.). Association for Supervision and Curriculum Development.

Winkelmes, M. A. (2013). Transparency in Teaching: Faculty Share Data and Improve Students' Learning. Liberal Education, 99(2), n2.